arXiv:1510.05284v1 [stat.ME] 18 Oct 2015

Privacy sets for constrained space-filling

Eva Benková
eva.benkova@jku.at
Johannes-Kepler-University Linz and Comenius University Bratislava
Radoslav Harman
radoslav.harman@fmph.uniba.sk
Comenius University Bratislava
Werner G. Müller
corresponding author: werner.mueller@jku.at
Johannes-Kepler-University Linz

October 20, 2015

Abstract

The paper provides typology for space filling into what we call "soft" and "hard" methods along with introducing the central notion of privacy sets for dealing with the latter. A heuristic algorithm based on this notion is presented and we compare its performance on some well-known examples.

1 Introduction

Most of the literature on space-filling designs attempts to achieve its aim by optimizing a prescribed objective measuring a degree of space-fillingness, such as maximin, minimax, etc., sometimes combined with an estimation or prediction oriented criterion (like suggested in [7]). Let us label those as "soft" space-filling methods. In contrast, "hard" space-filling methods ensure desirable properties by enforcing constraints on the designs, such that a secondary criterion can be used for optimization (e.g. D-optimality in the case of the Bridge-designs

of [3]). In this paper we intend to propose a general framework for the latter methods based on the central notion of so-called privacy sets, which allows us elegant and efficient formulations of design algorithms. Eventually we will compare our "hard" techniques to the more established "soft" procedures on several examples.

Formally, let $\mathcal{X} \subseteq \mathbb{R}^d$ be a non-empty design space and let ξ be a finite subset of $\mathcal{X}$ representing an experimental design, with each point of ξ being a design point of a single trial of the corresponding experiment. The notion of a "set" automatically implies that the order of design points is irrelevant, as well as that the replicated observations in the same design point are not allowed. Without loss of generality, the standard problem of optimal designs of experiments (cf. eg. [8]) is to maximize some criterion in the set of all permissible designs Γ, that is, to find $\xi^* \in \operatorname{argmax}_{\xi \in \Gamma}\{\Phi(\xi)\}$, where $\Gamma \subseteq \Xi$ and Ξ denotes the set of all designs. Here, the soft methods of space-filling focus on proposing and tuning the criterion Φ and then optimize on the set $\Gamma = \Xi$. Hard methods, in contrast, pose restrictions on the design, such that all designs from Γ satisfy a required minimum level of space-fillingness.

The statistical quality of one design compared to another can be assessed by calculating their mutual efficiency, which is defined as $\operatorname{eff}(\xi|\eta) = \Phi(\xi)/\Phi(\eta)$ for any $\xi, \eta \in \Xi$, $\Phi(\eta) > 0$. This quantity is particularly meaningful when the criterion Φ is positively homogeneous, which means $\Phi(\alpha\xi) = \alpha\Phi(\xi)$ for any $\xi \in \Xi$ and any $\alpha \geq 0$.

2 Privacy sets algorithm

Let us now define the central notion for our approach to the generation of (constrained) space-filling designs.

Definition 1. *For each $x \in \mathcal{X}$, let $\mathcal{P}(x) \subseteq \mathcal{X}$, $x \in \mathcal{P}(x)$, be a given privacy set of the point x. For any design $\xi \in \Xi$, let $\mathcal{P}(\xi) = \cup_{x \in \xi}\mathcal{P}(x)$ be the privacy set of the design ξ.*

We assume that there is a given upper limit on the size of the experiment $N \in \mathbb{N}$. A design ξ will then be called permissible, if $|\xi| \leq N$ and $x \notin \mathcal{P}(y)$ for all $x, y \in \xi$, $x \neq y$. A permissible design ξ will be called maximal permissible if it cannot be augmented without violation of some of the constraints, i.e., if $\xi \cup \{x\}$ is not permissible for all $x \in \mathcal{X} \setminus \xi$.

Assumption 1. *We assume that $|\xi| = N$ for any maximal permissible design.*

We would like to emphasize that the idea of privacy sets is not artificial, but can be used for example to ensure various space-filling properties. In fact, many of the widely-used designs, starting from the classical exact designs (cf. eg. [1]) restricted by at most one trial at each point, to popular Latin hypercube designs (LHD, see [6]), to Bridge designs (BD, see [3]), can be formulated in the terms of privacy sets. More specifically, we can set $\mathcal{P}(x) = \{x\}$ in the case of classical designs. For Latin hypercube designs the design space X is usually a finite grid and we have $\mathcal{P}(x) = \{y : \exists i \in \{1, \ldots, d\} : x_i = y_i\}$, where $x, y \in \mathbb{R}^d$ and x_i, y_i denote theirs i-th coordinates. Privacy sets for Bridge designs are given by equation (2) in Section 3.

Note that the use of privacy sets is meaningful not only for computer, but also for physical experiments. It covers, for example, time-separation constraints, where the designs space represents time, and consecutive trials must be performed at least δ time units apart (see, e.g., the second example in Section 5 of the paper [10]). In this case, $\mathcal{X} = \mathbb{R}_0^+$ and $\mathcal{P}(x) = (x - \delta, x + \delta) \cap \mathcal{X}$ for all $x \in \mathcal{X}$.

The set $\mathcal{P}(x)$ is typically some kind of a neighbourhood of x (containing also x itself) securing some "privacy" for it, although this is not strictly required by the definition itself. Using the privacy sets defined above we obtain the optimization problem

$$\xi^* \in \operatorname*{argmax}_{\xi \in \Xi} \{\Phi(\xi); |\xi| \leq N \text{ and } x \notin \mathcal{P}(y) \text{ for all } x, y \in \xi, x \neq y\}. \quad (1)$$

In the following, we present a general framework of exchange-type algorithm - Privacy Sets Algorithm (PSA) - for solving optimization problem (1). In general, the specification of individual steps depends greatly on the design space $\mathcal{X}$ and on the constraints given by the sets $\mathcal{P}(x)$, $x \in \mathcal{X}$, as well as on the optimization criterion Φ.

A characteristic feature of PSA is the ability to temporarily violate "privacy" of one or more design points. This offers a wider range of possibilities than when performing only permissible changes and prevents the algorithm from getting stuck, for instance when the privacy constraints are very strict. Let $\mathcal{A}(\xi) \subseteq \mathcal{X} \setminus \xi$ denote a set of "candidate points" that can possibly augment a maximal design ξ. The set $\mathcal{A}(\xi)$, in contrast to $\mathcal{P}(\xi)$, is not an attribute of the problem itself, but can be adjusted in order to ensure the optimum performance of the algorithm. Note that we do not require $\mathcal{A}(\xi)$ to contain solely permissible points $x \notin \mathcal{P}(\xi)$.

PSA does not pose any restrictions on the design space $\mathcal{X}$. Due to implementation reasons, however, we always assume a finite design

space of size $n \in \mathbb{N}$. If n is relatively small (up to thousands of design points, say), the implementation of PSA is rather simple for all kinds of privacy sets. This is mainly true because of the ability to store information about availability of each individual point of the design space. However, with n increasing, it becomes computationally intensive or even unfeasible to keep n-dimensional vectors in the computer memory and certain specific features of a particular class of privacy sets have to be considered.

One of the key parts of PSA is the efficient augmentation of a design that is not maximal with the remaining runs to achieve the full size N. This is done by employing the following forward-type procedure (Algorithm 1) which adds permissible design points one-by-one until a maximal design is obtained.

Algorithm 1: Greedy Procedure (GrP)

Input : A permissible design ξ, $|\xi| = N^* < N$.
Output: A permissible design ξ, $|\xi| = N$

1 **for** $i = 1 : N - N^*$ **do**
2 | Augment ξ with the point x, which maximizes $\Phi(\xi \cup \{x\})$ subject to $x \notin P(\xi)$.
3 **end**
4 **return** ξ

One-point permissible augmentation from Step 2 can be crucial in effective implementing of PSA algorithm. One of the straightforward solutions is to use the exhaustive enumeration of $\mathcal{X} \setminus \mathcal{P}(\xi)$ (for smaller problems) or to use a blind random search, that is, to choose the best point from a set of candidates sampled independently from $\mathcal{X} \setminus \mathcal{P}(\xi)$. This can be performed by a direct rejection method, which, however, tends to be very inefficient for some cases. Therefore, we recommend exploiting particularities of the privacy constraints, if possible (as for example in Section 3).

For any maximal design ξ and any point $x \in \mathcal{A}(\xi)$ let $\eta(\xi, x)$ be a possibly random maximal permissible design based on ξ, containing x. We can view $\{\eta(\xi, x) : x \in \mathcal{A}(\xi)\}$ as a randomly generated neighbourhood of the design ξ in the set of all maximal permissible designs, or, in other words, a set of slight "mutations" of the design ξ. The mutation procedure given by Algorithm 2 calculates $\eta(\xi, x)$ for any permissible design ξ and $x \in \mathcal{A}(\xi)$.

Algorithm 2: Mutation Procedure (MuP)

Input : A permissible design ξ, $|\xi| = N$ and a candidate point $x \in \mathcal{A}(\xi)$.

Output: A permissible design ξ, $|\xi| = N$

Remove from ξ all those points that belong in $\mathcal{P}(x)$.

Let $\xi = \xi \cup \{x\}$.

if $|\xi| = N + 1$ **then**

  Remove the design point from ξ that leads to the smallest drop in the criterion value.

else if $|\xi| < N$ **then**

  Augment the design ξ using GrP to the maximal design.

end

return ξ

The main body of the Privacy Sets Algorithm can then be written in the scheme of Algorithm (3).

Algorithm 3: Privacy Sets Algorithm (PSA)

Construct an initial design ξ using GrP.

repeat

  Set $\xi_{old} \leftarrow \xi$.

  Construct candidate set $\mathcal{A}(\xi)$.

  for $x \in \mathcal{A}(\xi)$ **do**

    Set $\eta \leftarrow \eta(\xi, x)$ using MuP.

    if $\Phi(\eta) > \Phi(\xi)$ **then**

      Set $\xi \leftarrow \eta$.

      break the for cycle

    end

  end

until $\Phi(\xi_{old}) \geq \Phi(\xi)$;

return ξ

3 Privacy sets algorithm for Bridge designs

In this section, we provide a particular version of PSA, which deals with so-called Bridge constraints. This term originates from [3], but we employ it for the constraints themselves, independent of the optimization criterion.

Let δ be a given positive constant. A design ξ will be called Bridge design, if ξ satisfies constraints given by the privacy sets defined for any $x \in \mathcal{X}$ by

$$\mathcal{P}(x) = \{y \in \mathcal{X} : |x_i - y_i| < \delta \text{ for some } i \in \{1, \dots, d\}\}. \tag{2}$$

Definition (2) suggests that the focus of Bridge constraints is on the space-fillingness of one-dimensional projections of the design points, since no two levels of a given factor can be closer than δ. Motivation for this requirement can be drawn from the non-collapsingness of the design in the case that some of the factors turn out to be irrelevant.

In this section, we assume $\mathcal{X} = [-1, 1]^d$, that is, every factor takes values in a bounded interval, which can be scaled to $[-1, 1]$. We assume that $\mathcal{X}$ is discretized into a grid of $n = L^d, L \in \mathbb{N}$, equally spaced points with each coordinate from $\left\{-1 + \frac{2k}{L-1}, k = 0, 1, 2, \dots, L-1\right\}$. Without loss of generality, we will choose the minimum spacing δ from the set $\left\{\frac{2k}{L-1}, k = 1, 2, \dots\right\}$.

The requirement of N experimental runs implies $\delta \leq 2/(N-1)$ and $L \geq N$. Note that for the special case $L = N$, we obtain the well-known Latin Hypercube designs (LHD). If $L > N$ and $\delta = 2/(L-1)$, the resulting design can be viewed as an "incomplete" LHD.

In general, the most computationally difficult part of PSA is the one-point permissible augmentation in Algorithm 1. Using the specific nature of Bridge designs defined on $[-1, 1]^d$, we implemented this step in two parts.

In the first part, we perform a blind random search by repeated sampling from $\mathcal{X} \setminus \mathcal{P}(\xi)$ and selecting that point which leads to the biggest increase of the criterion value. In the case of Bridge designs, it is enough to store just an $L \times d$ logical matrix representing permissible levels of factors. Points $x \in \mathcal{X} \setminus \mathcal{P}(\xi)$ can then be easily selected independently, coordinate by coordinate. In the case where additional

constraints on the design space are present (see Section 5), we used the rejection method.

In the second part, we tune the best point found by using a local search optimization procedure. Its main idea is to sequentially improve the position of the design point added in the first part, always varying only one coordinate at a time. Since all other design points remain unchanged, we are allowed to move only in the permissible area, where no collision with another design point occurs. The process of checking feasibility of the prospective design space point can be handled easily when considering Bridge restrictions (see the reasoning above). Note that this procedure does not require storing all design points (or regressors associated with the design points) in the computer memory. Therefore PSA for Bridge designs can be applied to very large design spaces.

4 Examples: D-optimal Bridge designs on a cubical design space

This section provides examples of Bridge designs for specific choices of the optimization criterion and the design space. We compare these to the results from [3], where the same settings were considered.

Let $\mathbf{f}(x)$ be an m-dimensional function at design point $x \in \xi$, $m \in \mathbb{N}$, usually representing a regression function in a linear model with uncorrelated homoscedastic errors.

In optimal experimental design, the most common objective function Φ is the D-criterion given by

$$\Phi_D(\xi) = \det(\mathbf{M}(\xi))^{1/m}, \tag{3}$$

where $\mathbf{M}(\xi) = \frac{1}{N}\sum_{x\in\xi} \mathbf{f}(x)\mathbf{f}^\top(x)$ is the standardized information matrix of the size $m \times m$. This definition of D-optimality ensures its positive homogeneity discussed in Section 1. If $\mathbf{f}(x)$ is a regressor in a linear model, a D-optimal design minimizes the generalized variance of the best linear unbiased estimator of the model parameter, cf. [8].

This criterion leads to the optimization problem

$$\xi^* \in \operatorname*{argmax}_{\xi\in\Xi}\{\Phi_D(\xi); |\xi| \leq N \text{ and } |x_i - y_i| \geq \delta \text{ for all } i = 1, \ldots, d\}.$$

It is natural to require the design space to be a d-dimensional hypercube $[-1, 1]^d$ (for example when conducting a computer experiment), hence we restrict ourselves to this assumption.

With the number of trials N given, the density of the design space grid needs to be chosen. First, we know that the minimal distance δ cannot exceed $2/(N-1)$, but it is recommended to set it to a smaller value in order to maintain a certain level of "freedom". Now, with the value of δ determined, we can set $L = \lfloor 2k/\delta \rfloor + 1$, where $k \in \mathbb{N}$ and $\lfloor . \rfloor$ represents the lower integer part. In [3] the parameter k is always set to 1, yielding the "incomplete" LHD defined in the previous section. In general, the choice of small L can accelerate the algorithm, but also impair the quality of the resulting design, which suggests that a certain compromise should be made.

In the following examples, we illustrate the performance of our algorithm applied to the problems of various dimensions. We compare the results obtained by our algorithm to the results from [3], based on the relative efficiency of the best designs found in a fixed time interval.

To make the comparison as fair as possible, we implemented both algorithms in Matlab computing environment. The algorithm of [3] was translated from JMP scripting language which was available on the supplement area of the published article. We used 64 bit Windows 7 system running an Intel Core i3-4000M CPU processor at 2.40 GHz with 4 GB of RAM.

4.1 Bridge designs for 2 factors in 21 runs

As the first example, we consider the two-dimensional design space with 21 trials to be allocated. Since this problem instance is rather small, we set the computing time to 60 seconds and run both algorithms several times, until the whole time given is spent. The designs with the highest criterion value found by PSA are displayed in Figure 1, with the minimum spacing constant δ set to 0.05 and 0.025 for both the linear regression function $\mathbf{f}(x) = (1, x_1, x_2)^\top$ and the full quadratic regression function $\mathbf{f}(x) = (1, x_1, x_2, x_1^2, x_2^2, x_1 x_2)^\top$.

When compared to the results of [3], we observe better results of our algorithm in all four examined situations. The mutual efficiencies of the designs of [3] relative to our designs presented in the figures 1(a), 1(b), 1(c), 1(d) are equal to 0.79, 0.82, 0.96, and 0.96, in the respective order.

This suggests that our algorithm provides better results especially when the spacing constraint δ is rather large, which is not surprising. First, the finer design space grid automatically makes the relative difference in the output designs less significant. Second, the greater

the value of δ is in comparison to its upper bound $2/(N-1)$, the less space there is for "manoeuvering" for the algorithm. For example, the marginal case $\delta = 2/(N-1)$ leads to the standard LHD, where no observation can be added without violating some of the privacy sets constraints. These restrictions would fully disable the coordinate-exchanges in the algorithm of [3], but could still be handled by PSA.

4.2 A numerical study on Bridge designs

The comparison of the algorithm of [3] and PSA can be extended into more-dimensional cases as well. We performed a small comparative numerical study on a few examples of quadratic regression of various dimensions, similar to the examples provided in Sections 3 and 4 of [3]. For every example, we ran the algorithms for a restricted time, observing the time dependence of the criterion value of the actually best design found by the algorithm. We repeated this procedure several times in order to provide responses from multiple random starts.

In Figure 2, we present results of the two competing algorithms for dimensions $d = 2, 4, 6, 8$ with the numbers of trials $N = 21, 41, 61, 81$, respectively. The minimum spacing δ was in all four cases set to the value $\delta = 1/(N-1)$, which corresponds to the value recommended in [3]. Every t seconds, we plotted the criterion values of the best designs found by the algorithm of [3] (represented by the red lines) and PSA (represented by the blue lines) versus the time displayed on the x-axis. Both algorithms were restarted 5 times, yielding 5 red and 5 blue lines for each problem instance.

The total computational time T, as well as the time period t, were chosen such that they increase with the increasing size of the problem. If an algorithm terminated during the given time T, it was automatically restarted and the resulting value found by its run was stored in the memory. These restarts are denoted by the red and the blue diamonds. We considered the D-optimality criterion given in (3) and plotted its values on the y-axis.

Note that PSA was able to significantly outperform the algorithm of [3] in all four examined cases. We also note that PSA yielded an efficient design in a relatively short time, which suggests that it can be reasonable to stop PSA before reaching an actual local optimum and reduce the execution time.

Although the D-criterion is of central interest in this section, one could be possibly interested in space-filling properties of the resulting designs, as well. There exist plenty of space-filling criteria to choose from when comparing various designs (e.g., recall that Bridge constraints themselves force a certain one-dimensional space-fillingness).

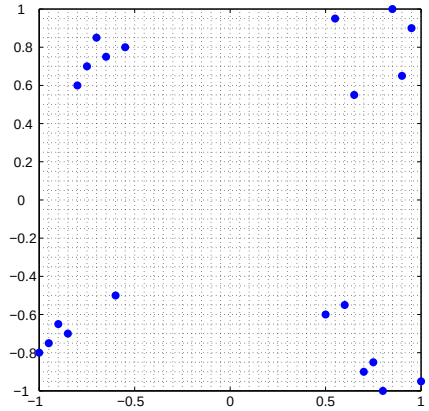

(a) linear regressor, $\delta = 0.05$

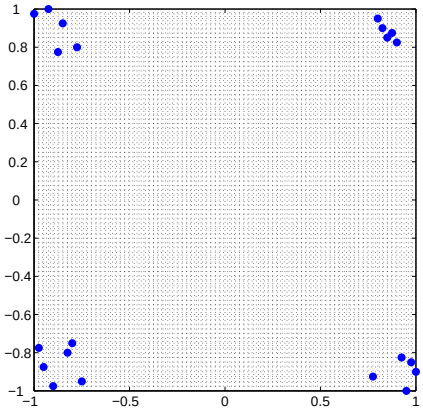

(b) linear regressor, $\delta = 0.025$

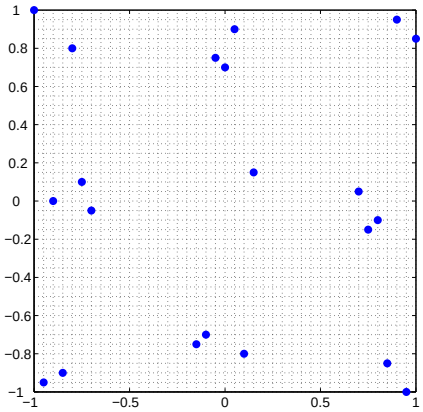

(c) full quadratic regressor, $\delta = 0.05$

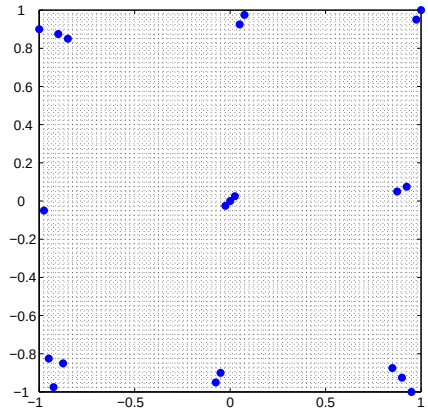

(d) full quadratic regressor, $\delta = 0.025$

Figure 1: Bridge designs for $N = 21$ and $d = 2$ found in 60 sec. In 1(a) and 1(b), the linear regressor was used and the minimal distance δ was set to 0.05 (1(a)) and 0.025 (1(b)). In 1(c) and 1(d) the full quadratic regressor was used and the minimal distance δ was set to 0.05 (1(c)) and 0.025 (1(d)). Efficiencies of the best designs found by the algorithm of Jones et al. relative to the designs presented in the figures 1(a), 1(c),1(b),1(d) are 0.79, 0.82, 0.96 and 0.96, respectively.

We choose one of them, provided in [3], which is based on the distances of the selected points to the nearest design point.

In Figure 3 we study 4 different 32-point designs in 6 factors. For every design, we display a box plot of distances to the nearest design point of 64 vertices and 10000 points randomly sampled from $[-1, 1]^d$.

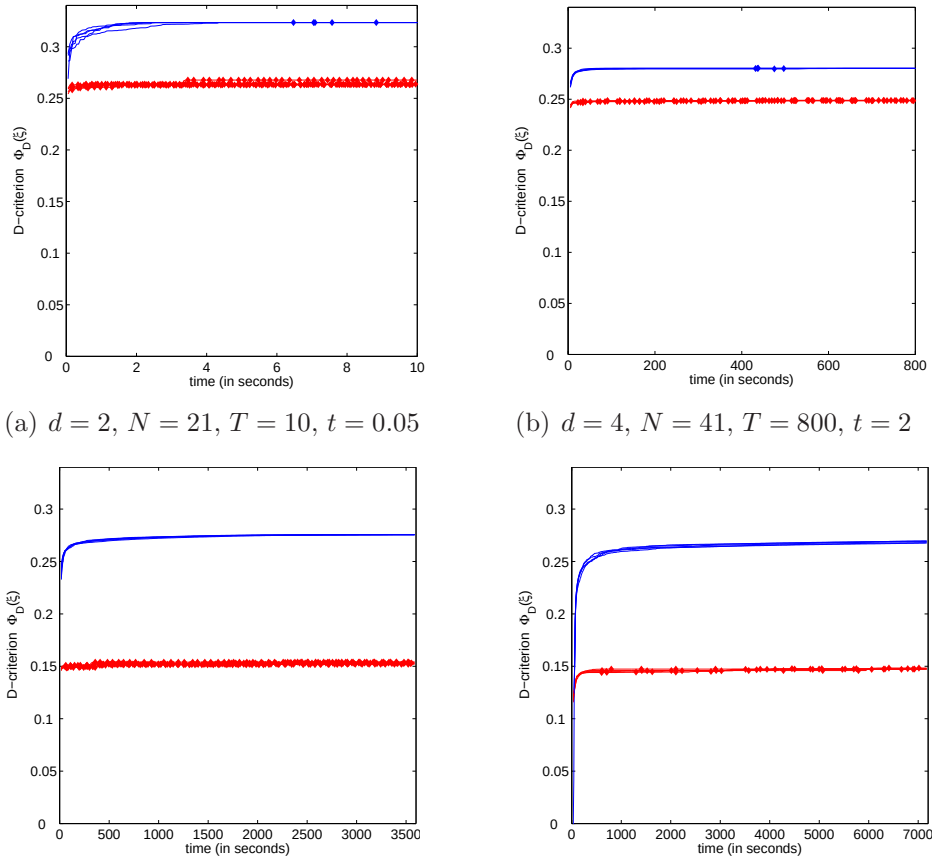

(a) $d = 2$, $N = 21$, $T = 10$, $t = 0.05$

(b) $d = 4$, $N = 41$, $T = 800$, $t = 2$

(c) $d = 6$, $N = 61$, $T = 3600$, $t = 5$

(d) $d = 8$, $N = 81$, $T = 7200$, $t = 10$

Figure 2: Comparison of the performance of the algorithm of [3] (red lines) and PSA (blue lines). Y-axis shows the best D-optimality values found by the time displayed on the x-axis (in seconds). The red and the blue diamonds denote restarts of the corresponding algorithms. For every example both algorithms were ran 5 times for the time period T.

The designs compared are gained either by the algorithm of [3](“J”) or by Privacy sets algorithm (“PSA”), with the minimum spacing δ set to either $\delta^* = 1/(N-1) = 1/31$ or $\delta^*/2 = 1/62$.

The lower part of Figure 3 shows D-optimality values of the displayed designs, denoted by the black dots. Roughly speaking, we can say that the behaviour of these values approximately follows the behaviour of the corresponding box plots. In other words, this means that the design which is the best according to the distances to the

nearest design point ("J δ^*"), has the worst value of D-criterion, and vice versa. This is not surprising, since D-optimal designs have the tendency to concentrate in a few points of the design space and do not possess any particular space-filling properties.

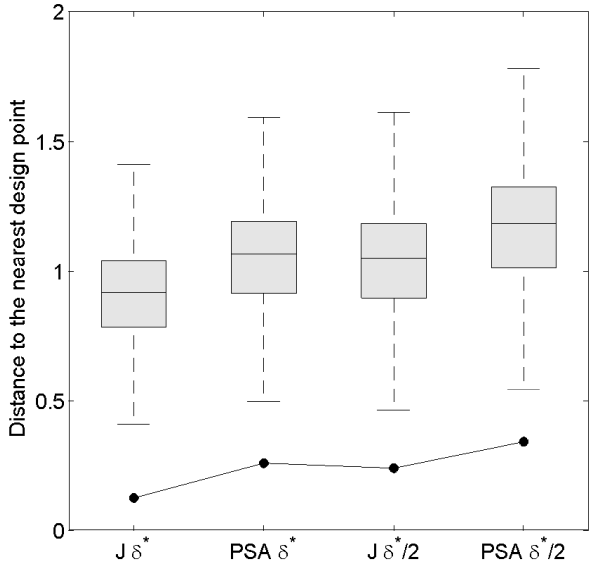

Figure 3: Box plots of distances of 10000 randomly selected points and 64 vertices to the nearest design points for $d = 6$ and $N = 32$. The presented designs are divided according to the minimal distance (δ^* or $\delta^*/2$) and the algorithm used ([3] or PSA). The black dots in the bottom part represent the D-criterion values. Although the results of [3] are worse in terms of D-optimality, they are slightly better in terms of space-fillingness.

5 Space-filling designs on a constrained design space

Note that we can practically think of any space-filling design as a particular instance of a 'Bridge design' and that this does not necessarily include D-optimality and a cubical design region. In fact, for any such design, it is enough to satisfy (2), that is, restrictions ensuring non-collapsing properties of the design. As an example, we present in this section space-filling on a constrained design space.

For that purpose, we choose one of the space-filling criteria to be optimized, which means that we combine together 'soft' and 'hard' methods described in Section 1. We consider the average reciprocal distance (ARD) criterion, modified such that the optimal design has good projection properties onto a given set of subspaces of the design space (as introduced in [2]). Let $J \subset \{1, 2, \ldots, d\}$ be a nonempty index set of dimensions of subspaces we would like to consider and let $\mathcal{X}_j$ denote the set of all $\binom{d}{j}$ standard coordinate subspaces of dimension d for every $j \in J$.

The idea of the ARD criterion is that the average reciprocal pairwise distances of design points should be minimized. Hence, in order to remain consistent with the 'maximization policy' adopted in this paper (see equation (1)), we use the following formulation of ARD:

$$\Phi_{ARD}(\xi) = \left(\frac{1}{\binom{N}{2} \sum_{j \in J} \binom{d}{j}} \sum_{j \in J} \sum_{\mathcal{Y} \in \mathcal{X}_j} \sum_{\substack{x,y \in \xi \\ x \neq y}} \left(\frac{j^{1/z}}{\rho_z(x^*_{\mathcal{Y}}, y^*_{\mathcal{Y}})} \right)^{\lambda} \right)^{-1/\lambda}, \quad (4)$$

where $z \geq 1$ and $\lambda \geq 1$ are given constants, $x^*_{\mathcal{Y}}$ is the projection of $x \in \mathcal{X}$ onto subspace $\mathcal{Y}$, and ρ_z is for any couple $x = (x_1, \ldots, x_d)^\top$, $y = (y_1, \ldots, y_d)^\top$ defined by

$$\rho_z(x, y) = \left(\sum_{i=1}^{d} |x_i - y_i|^z \right)^{1/z}.$$

Each design point has to be selected from a design space, which in this case will be some linearly constrained region. However, the PSA algorithm for Bridge designs on cubical regions, described in Section 3, can be rather straightforwardly adapted for the constrained design regions.

Without loss of generality, assume that the design space $\mathcal{X}$ is a subset of $[-1, 1]^d$ with some additional linear constraints $Ax \leq b$ to be satisfied for every $x \in \mathcal{X}$, where A is an $k \times d$ matrix , $k \in \mathbb{N}$, and $b \in \mathbb{R}^d$. For a given number of trials N, let us have the design space discretized by first making a grid of L^d points on $[-1, 1]^d$ (see Section 3), and then accepting only those satisfying $Ax \leq b$. The parameter L, determining the density of the grid, has to be chosen large enough, such that not only a permissible design $\xi \in \mathcal{X}$ exists, but also that Assumption 1 is satisfied.

The only question is how to implement the one-point permissible augmentation from Algorithm 1. We utilize blind random search on the set $\mathcal{X} \setminus \mathcal{P}(\xi)$ and select the best point found. In the case of a Bridge design on $[-1,1]^d$, we have a simple tool to sample from $\mathcal{X} \setminus \mathcal{P}(\xi)$, by keeping track of permissible and non-permissible levels of individual factors, see Section 3. If additional linear restrictions are present, $x \in \mathcal{X} \setminus \mathcal{P}(\xi)$ can be generated in the same way and then simply accepted if the condition $Ax \leq b$ holds and rejected otherwise. Clearly, effectiveness of this rejection method depends on the design region defined by the constraints, as well as on the dimension d. In every step, we have to check the restriction on the design space, but we do not have to check the collision with other design points in terms of privacy sets (due to the convenient Bridge constraints), which makes the rejection method more efficient than in the case of general privacy sets.

Figure 4 displays three resulting designs obtained by PSA for different variants of the ARD criterion. The design space is in all three situations the square $[-1,1]^2$ additionally restricted by $\frac{1}{2}x_1 - x_2 \leq \frac{1}{2}$. We considered designs with $N = 100$ runs and the minimum spacing parameter $\delta = \frac{2}{120-1}$. We note that in this case, it is not possible to set $\delta = \frac{2}{N-1}$ ("LHD-setting"), since Assumption 1 would not be fulfilled and the PSA algorithm could not be executed.

The parameters of the ARD criterion (5) were set to $z = 1$ and $\lambda = 1$. The nonempty index set $J \subseteq \{1,2\}$ varies in figures 4(a), 4(b), 4(c) through all three possibilities, leading to different output designs. Space-filling properties of these designs for various settings of ARD criterion are compared in Table 1.

Design	Fig. 4(a)	Fig. 4(b)	Fig. 4(c)
Criterion			
ARD $J = \{1\}$	4.2497	4.4106	4.2718
ARD $J = \{2\}$	2.6351	2.1876	2.2473
ARD $J = \{1,2\}$	3.7115	3.6696	3.5970

Table 1: Space-filling properties of designs displayed in Figure 4 measured by ARD criterion for different values of $J \subseteq \{1,2\}$.

Histograms in Figure 4, as well as the first row of Table 1, measure space-fillingness of one-dimensional projections of the designs. Ideally,

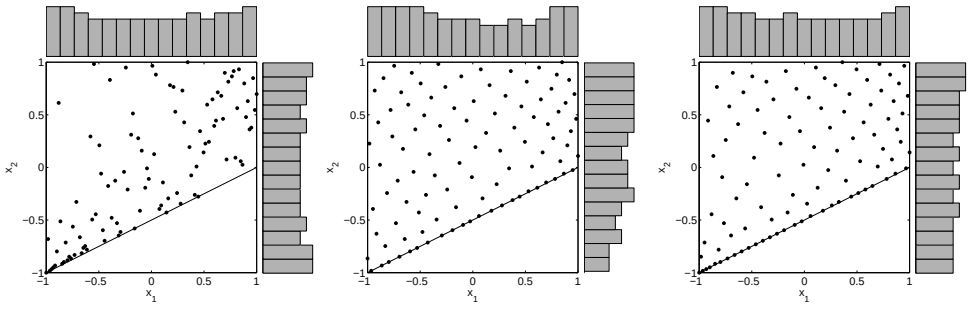

(a) $J = \{1\}$ (b) $J = \{2\}$ (c) $J = \{1, 2\}$

Figure 4: Designs on the two-dimensional constrained design region, obtained by PSA optimizing the ARD criterion for $z = 1$, $\lambda = 1$ and different values of $J \subseteq \{1, 2\}$. PSA allocated 100 trials for the minimum spacing constant $\delta = \frac{2}{120-1}$.

we would like to have these projections as uniformly distributed as possible (corresponding to Latin Hypercube constraints), which would ensure non-collapsing properties in one-dimension. One could quite easily think of a heuristics forcing perfectly uniform one-dimensional projections of a design. However, we think it can be more beneficial to "relax" constraints from LHD to BD for some smaller value of δ, if this allows us to improve another optimization criterion, e.g. the criterion assessing space-filling properties of more-dimensional projections (ARD for $J = \{2\}$ or $J = \{1, 2\}$). In this sense, we can compare designs of Figure 4 to the design from the paper [9], given in Figure 4 (a) of Section 2.3, which differs only in the scaling of the design space. Although the design of [9] strictly satisfies Latin Hypercube constraints (see the histograms of one-dimensional projections), it completely ignores space-fillingness in other dimensions.

Figure 5 presents an example of a design resulting from PSA for ARD criterion in three factors. The three dimensional design space cube $[-1, 1]^3$ is additionally constrained by $\frac{2}{3}x_1 - x_2 \leq \frac{1}{3}$ and $\frac{3}{4}x_2 - x_3 \leq \frac{1}{4}$. The design consists of $N = 100$ trials with their one-dimensional projections not closer together than $\delta = \frac{2}{140-1}$. The ARD criterion was employed for $z = 1$, $\lambda = 1$ and $J = \{2, 3\}$, which means that we combined the 'hard' method forcing one-dimensional space-fillingness and the 'soft' method ensuring space-filling properties for dimensions 2 and 3.

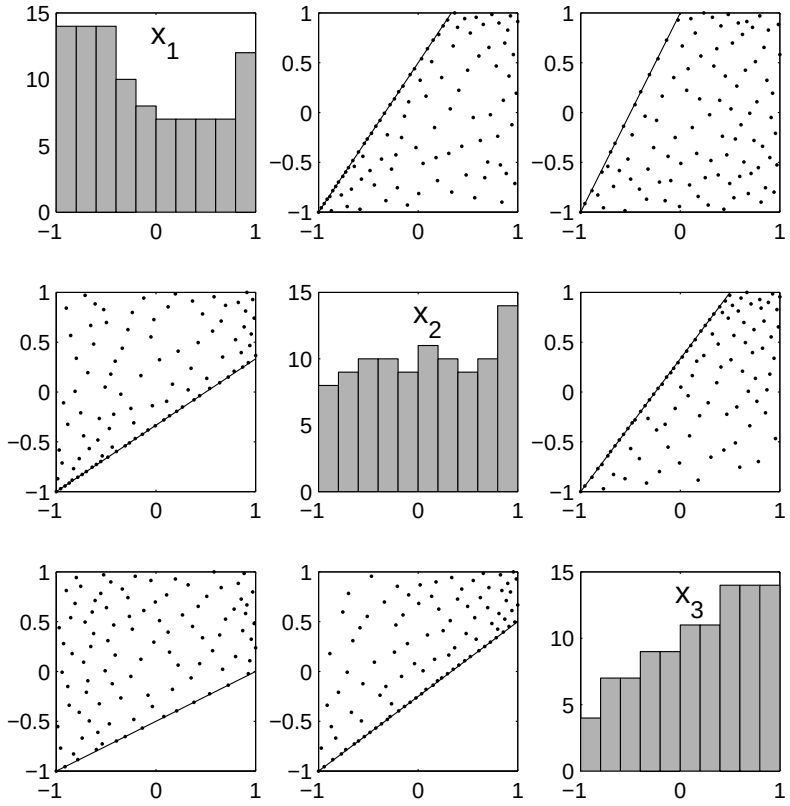

Figure 5: Design on the three-dimensional constrained design region, obtained by PSA optimizing the ARD criterion for $z = 1$, $\lambda = 1$ and $J = \{2, 3\}$. PSA allocated 100 trials for the minimum spacing constant $\delta = \frac{2}{140-1}$.

6 Conclusions

As we believe to have shown in this paper the notion of privacy sets is central to the understanding and interpretation of "hard" space-filling. The algorithms based on this notion are extremely flexible and can be used in a great variety of situations. As we have also shown we are not restricted to the "hard" space-filling paradigm, but we can encompass "soft" methods and combinations as well.

This is perhaps emphasized best by relating to the recently proposed MaxPro criterion of [5], which was designed for balancing out

performance in all dimensions and is given by

$$\Phi_{MaxPro}(\xi) = \sum_{\substack{x,y\in\xi \\ x\neq y}} \left(\frac{1}{\Pi_{i=1}^{d} \|x_i^* - y_i^*\|^z} \right), \tag{5}$$

where usually $z = 2$. In the discussion of [4] we have demonstrated that our techniques can be easily adapted for these purposes. These and other useful extension will be a matter of our further research.

Acknowledgements

The research of the first author has been financially supported by the ANR/FWF grant I-833-N18, the research of the second author was supported by the VEGA 1/0163/13 grant of the Slovak Scientific Grant Agency.

References

[1] George Casella. *Statistical Design*. Springer New York, New York, NY, 2008.

[2] Danel Draguljic, Angela M. Dean, and Thomas J. Santner. Non-collapsing space-filling designs for bounded nonrectangular regions. *Technometrics*, 54(2):169–178, 2012.

[3] Bradley Jones, Rachel T. Silvestrini, Douglas C. Montgomery, and David M. Steinberg. Bridge Designs for Modeling Systems with Low Noise. *Technometrics*, 57(2):155–163, 2014.

[4] V. Roshan Joseph. Space-filling Designs for Computer-Experiments: A Review. *Quality Engineering*, forthcoming, 2015.

[5] V. Roshan Joseph, Evren Gul, and Shan Ba. Maximum projection designs for computer experiments. *Biometrika*, 102(2):371–380, 2015.

[6] M. D. McKay, R. J. Beckman, and W. J. Conover. A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code. *Technometrics*, 21(2):239–245, 1979.

[7] Max D. Morris and Toby J. Mitchell. Exploratory designs for computational experiments. *Journal of Statistical Planning and Inference*, 43(3):381–402, 1995.

[8] Andrej Pázman. *Foundations of Optimum Experimental Design (Mathematics and its Applications)*. Reidel Publ. Comp., Dodrecht, 1986.

[9] Matthieu Petelet, Bertrand Iooss, Olivier Asserin, and Alexandre Loredo. Latin hypercube sampling with inequality constraints. *AStA Advances in Statistical Analysis*, 94(4):325–339, 2010.

[10] Guillaume Sagnol and Radoslav Harman. Computing exact d-optimal designs by mixed integer second-order cone programming. *Ann. Statist.*, 43(5):2198–2224, 2015.